

02/25/00



02-28-00

A

JC525 U.S. PTO  
09/513706  
02/25/00

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

**CERTIFICATE OF EXPRESS MAILING**

I hereby certify that this paper and the documents and/or fees referred to as attached therein are being deposited with the United States Postal Service on February 25, 2000 in an envelope as "Express Mail Post Office to Addressee" service under 37 CFR §1.10, Mailing Label Number **EL117278326US**, addressed to the Assistant Commissioner for Patents, Washington, DC 20231.

Attorney Docket No.: SUN1P398/P4686

First Named Inventor: LODRIGE

*Diana Ellis*

Diana Ellis

**UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR. § 1.53(b))**

Assistant Commissioner for Patents  
Box Patent Application  
Washington, DC 20231

☐ Duplicate for  
fee processing

**For:** This is a request for filing a patent application under 37 CFR. § 1.53(b) in the name of inventors:  
Paul F. Lodrige, Randy S. Fishel

**For:** METHOD AND APPARATUS FOR CONCURRENT PROPAGATION OF DATA BETWEEN  
SOFTWARE MODULES

Application Elements:

- ☒ 22 Pages of Specification, Claims and Abstract
- ☒ 08 Sheets of Formal Drawings
- ☒ 02 Pages Combined Declaration and Power of Attorney (Unexecuted)

Accompanying Application Parts:

- ☐ Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)
- ☐ Separate Power of Attorney by Assignee
- ☐ 37 CFR 3.73(b) Statement by Assignee
- ☐ Information Disclosure Statement with Form PTO-1449
  - ☐ Copies of IDS Citations
- ☐ Preliminary Amendment
- ☒ Return Receipt Postcard

**PLEASE DO NOT CHARGE FEES AT THIS TIME**

☒ The Commissioner is authorized to charge any fees beyond the amount enclosed which may be required, or to credit any overpayment, to Deposit Account No. 50-0388 (Order No. SUN1P398).

General Authorization for Petition for Extension of Time (37 CFR §1.136)

☒ Applicants hereby make and generally authorize any Petitions for Extensions of Time as may be needed for any subsequent filings. The Commissioner is also authorized to charge any extension fees under 37 CFR §1.17 as may be needed to Deposit Account No. 50-0388 (Order No. SUN1P398).

☒ Please send correspondence to the following address:

**Customer Number 022434**  
BEYER WEAVER THOMAS & NGUYEN, LLP  
P.O. Box 130  
Mountain View, CA 94042-0130  
Telephone (650) 961-8300  
Fax (650) 961-8301



Date:

2/25/2000

A handwritten signature in dark ink, appearing to read "C. Douglass Thomas", written over a horizontal line.

**C. Douglass Thomas**  
Registration No. 32,947

## **PATENT APPLICATION**

### **METHOD AND APPARATUS FOR CONCURRENT PROPAGATION OF DATA BETWEEN SOFTWARE MODULES**

Inventors:

1. Paul F. Lodrige  
784 Holly Oak Drive  
Palo Alto, CA 94303  
US Citizen
2. Randy S. Fishel  
854 Laburnum  
Sunnyvale, CA 94086  
US Citizen

Assignee:

Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303

BEYER WEAVER THOMAS & NGUYEN, LLP  
P.O. Box 130  
Mountain View, CA 94042  
Telephone (650) 961-8300

# METHOD AND APPARATUS FOR CONCURRENT PROPAGATION OF DATA BETWEEN SOFTWARE MODULES

## CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to U.S. Patent Application No. \_\_\_\_\_  
(Att.Dkt.No. SUN1P397/P4611), entitled "METHOD AND APPARATUS FOR  
MANAGING DATA PROPAGATION BETWEEN SOFTWARE MODULES ", filed  
concurrently herewith, and hereby incorporated herein by reference.

## BACKGROUND OF THE INVENTION

### **1. Field of the Invention**

The present invention relates to computing systems and, more  
particularly, to data communications for computing systems.

### **2. Description of the Related Art**

Recent developments in data communication for computing systems have  
software modules in a layered model. One feature of UNIX System V uses a  
layered model of software modules, referred to herein as the STREAMS model.  
The STREAMS model provides a standard way of dynamically building and  
passing messages through software modules that are placed in layers in a  
protocol stack. In the STREAMS programming model, the protocol stack can be  
dynamically changed, e.g., software modules can be added or removed (pushed  
and popped) at run-time. Broadly speaking, a "stream" generally refers to an  
instance of a full-duplex path using the model and data communication facilities  
between a process in user space and a driver which is typically located within  
the kernel space of the computing system. In other words, a stream can be  
described as a data path that passes data in both directions between a stream  
driver in the kernel space and a process in user space.

Fig. 1 illustrates a STREAMS programming environment 100 including a stream head 102, stream modules 104 and 106, and a streams driver 108. An application program can create a "stream" by opening a device 110. Initially, the stream includes only the stream head 102 and stream driver 108. The stream head 102 is the end of the stream closet to the user process and serves as a user interface between the stream and the user process. Similarly, the stream driver 108 is the end of the stream closest to the device 110 and transfers data between the kernel and the device 110.

After the stream head 104 and stream driver 108 are provided, one or more stream modules, such as stream modules 104 and 106, can be pushed on the stream between the stream head 102 and stream driver 108. An application can dynamically add or remove (push or pop) stream modules on the stream stack at run-time. Each of the stream modules 104 and 106 includes a defined set of kernel-level routines and data structures to process data that passes through these modules. For example, stream modules can operate to convert lower case to upper case, or to add network routing information, etc.

As depicted in Fig. 1, each of the stream head 102, the stream modules 104 and 106, and the stream driver 108 has a pair of queues that can be used as containers for holding messages (e.g., data) traveling upstream or downstream over the stream. For example, a down-queue 112 and a up-queue 114 hold blocks of messages (data) for the stream module 104 and respectively pass the data up the stream to the stream head 102 or down the stream to the stream driver 108. The messages can be ordered within the queues 112 and 114 in first-in first-out basis (FIFO), perhaps according to assigned priorities.

It should be noted that in some situations messages cannot be placed in queues 112 and 114. For example, when a stream queue has reached it allotted size, messages can no longer be placed in that queue. As another example, messages cannot be placed in the queues 112 and 114 when other processing threads have acquired their software locks. In such cases, messages are stored in another queue that can serve as a back-up queue, herein referred to as a

"synchronization queue". For example, synchronization queues 116 and 118 depicted in Fig. 1, respectively hold messages that could not be placed into the queues 112 and 114. It should be noted that in the STREAMS model, some stream modules (e.g., Internet Protocol (IP) module) are confined to having only one synchronization queue. As a result, for some stream modules there is only one synchronization queue available to serve as a backup queue. For example, a synchronization queue 120 which is used to hold the messages for the second software module 106.

One problem with conventional implementations of the STREAMS model is that in some instances messages cannot be propagated concurrently by two or more threads (or processes) running on different processors. For example, when a stream module has only one synchronization queue available, access to the synchronization queue cannot conventionally be achieved by two or more threads. This is partly attributed to the fact that in conventional implementations of the STREAMS model that messages can be intermixed in one synchronization queue regardless of their type and/or relative priority. Accordingly, access to the synchronization queue is limited to only one thread at a time. Typically, in conventional implementations of the STREAMS model, a synchronization queue lock is provided to exclude all the threads from accessing the synchronization queue except one single thread which has acquired (or possesses) the synchronization queue lock. For example, a synchronization queue lock can be provided to limit access to the synchronization queue 120 to only one thread at any given time (not shown).

In view of the foregoing, there is a need for improved methods for managing data propagation between software modules.

## SUMMARY OF THE INVENTION

Broadly speaking, the invention relates to techniques for concurrent propagation of data through software modules used by computer systems. More particularly, the invention allows a plurality of threads running on different processors to propagate data between software modules by providing concurrent access to queued data by multiple threads. The invention is particularly well suited for use with the STREAMS model that uses software models arranged in a stack to provide data communications.

10 The invention can be implemented in numerous ways, including a system, an apparatus, a method or a computer readable medium. Several embodiments of the invention are discussed below.

As a method for managing flow of messages between two software modules, one embodiment of the invention includes that acts of: (a) determining whether a first message can be propagated by a first thread running on a first processor between the two software modules while allowing a second thread running on a second processor to propagate a second message between the two software modules; and (b) propagating the first message between the two software modules while allowing the second thread to propagate the second message between the two software modules when said determining (a) determines that the first message can be propagated by the first thread while allowing the second thread to propagate the second message between the two software modules.

As a computer system, one embodiment of the invention includes: a plurality of processors; a first and a second software modules with the second software module having a main queue suitable for storing messages and an auxiliary queue suitable for storing messages that are not stored in the main queue; and a propagation controller operating to enable at least two processors

of said plurality of processors to concurrently propagate messages between the first and the second software modules.

As a computer readable media including computer program code for managing flow of messages between two software modules, one embodiment of the invention includes: computer program code for determining whether a first message can be propagated by a first thread running on a first processor between the two software modules while allowing a second thread running on a second processor to propagate a second message between the two software modules; and computer program code for propagating the first message between the two software modules while allowing the second thread to propagate the second message between the two software modules when said computer program code for determining determines that the first message can be propagated by the first thread while allowing the second thread to propagate the second message between the two software modules.

The advantages of the invention are numerous. Different embodiments or implementations may have one or more of the following advantages. One advantage of the invention is that data can be propagated concurrently between software modules by a plurality of threads. Another advantage of the invention is that lock-free data propagation is made possible. Yet another advantage of the invention is that propagation of data can be achieved more quickly and more efficiently.

Other aspects and advantages of the invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.



## **BRIEF DESCRIPTION OF THE DRAWINGS**

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein  
5 like reference numerals designate like structural elements, and in which:

Fig. 1 illustrates a STREAMS programming environment including a stream head, two stream modules, and a streams driver.

Fig. 2 illustrates a computing environment including an enhanced propagation system suitable for allowing concurrent propagation of messages  
10 between software modules in accordance with one embodiment of the present invention.

Fig. 3 illustrates a data flow management method configured according to one embodiment of the present invention.

Figs. 4A and 4B illustrate a data flow management method configured  
15 according to another embodiment of the invention.

Figs. 5A –5D illustrate a computing environment in which a representative message is propagated from a (main) queue of a source software module to a destination software module 504 in accordance with various situations of the present invention.

20 Figs. 6A and 6B illustrate an example of a computer system that may be used in accordance with the invention.

## **DETAILED DESCRIPTION OF THE INVENTION**

The invention pertains to techniques for allowing multiple processes (or threads) to concurrently propagate data through software modules used by computer systems. More particularly, the invention allows a plurality of threads running on different processors to propagate data between software modules by providing concurrent access to queued data by multiple threads. The invention is particularly well suited for use with the STREAMS model that uses software models arranged in a stack to provide data communications.

Embodiments of the invention are discussed below with reference to Figs. 2 –6B. However, those skilled in the art will readily appreciate that the detailed description given herein with respect to these figures is for explanatory purposes as the invention extends beyond these limited embodiments.

Fig. 2 illustrates a computing environment 200 including an enhanced propagation system 202 suitable for allowing concurrent propagation of messages between first and second software modules 204 and 206 in accordance with one embodiment of the present invention. More particularly, the enhanced propagation system 202 allows multiple threads running on different processors to propagate data (e.g., messages) between a queue of the first software module 204 and a synchronization queue of the second software modules 206, or vice versa. The first and second software modules 204 and 206 can be part of a stream in a STREAMS model operating environment. As shown in Fig. 2, a first queue 208 and a synchronization queue 210 are provided for the first software module 204. In the STREAMS model operating environment, messages are processed by the first software module 204 and propagated to a subsequent software module (e.g., the second software module 206). Messages processed by the first software module 204 can be stored in the first queue 208. The first synchronization queue 210 can function as a back up queue for the first queue 208. Thus, the messages that cannot be placed in the first queue 208 can often be stored in the first synchronization queue 210. The second queue 212 and synchronization queue 214 store messages for the second software module

206, in the same manner. Although discussion focuses on the propagation of messages from the first software module 204 to the second software module 206, it should be understood that data propagation also likewise occurs from the second software module to the first software module 204.

5 As shown in Fig. 2, the enhanced propagation system 202 includes a propagation controller 216, a thread-counter 218, and a synchronization queue counter 220. The propagation controller 216 facilitates concurrent propagation of messages between the first and second software modules 204 and 206 by permitting concurrent access to the synchronization queue 210 and/or  
10 synchronization queue 214. The thread-counter 218 provides a thread-count for each processor (e.g., central processing unit) that is available in the computing environment 200 to propagate data. The thread-count for each processor indicates the number of threads running on that processor that are propagating messages between the first and second software modules 204 and 206. For  
15 example, a thread count of zero for a processor can indicate that no threads running on the processor are currently propagating data. In this way, the propagation controller 216 can determine whether a processor is available to propagate data. When a thread on a processor is assigned to propagate data, the appropriate thread-counter is updated (e.g., incremented to one). The  
20 synchronization queue counter 220 can provide information as to whether a synchronization queue is being accessed by a thread. For example, Count-1 and Count-2 of the synchronization queue counter 220 can be provided respectively for the first and second synchronization queues 210 and 214. The propagation controller 216 can use the synchronization queue counter 220 to determine  
25 whether synchronization queues 210 and 214 are being accessed by a thread. The operation of the enhanced propagation system 202 will be discussed in greater detail below.

Fig. 3 illustrates a data flow management method 300 configured according to one embodiment of the present invention. The data flow  
30 management method 300 can be utilized by the enhanced propagation system

202 of Fig. 2 to allow concurrent propagation of messages between two software modules by multiple threads running on different processors.

Initially, at operation 302, a determination is made as to whether concurrent propagation of data is possible. Generally, concurrent propagation of data is possible when there is no need to prevent other threads from propagating data while a particular thread is propagating data. If operation 302 determines that concurrent propagation of data is not possible, the data flow management method 300 proceeds to operation 304 where a data flow lock is acquired. For example, in the STREAMS model, the data flow lock (e.g., a streams-head lock) prevents any thread, except the thread that possess the lock, from propagating data in the stream. Next, at operation 306, a synchronization queue lock for the synchronization queue of the destination software module is acquired. For example, in the STREAMS model, the lock can be a streams synchronization queue lock which once acquired prevents any thread except the particular thread that possess the lock accessing the synchronization queue. Next, at operation 308, the message is propagated from the source software module (e.g., its queue) to the synchronization queue of destination software module. The data flow lock and synchronization queue lock are both released following propagation of messages at operation 310.

On the other hand, if the operation 302 determines that concurrent propagation of data is possible, the data flow management method 300 proceeds to operation 312 where the thread-count for the appropriate processor is updated (e.g. incremented from zero to one). The appropriate processor is the processor running the thread seeking to propagate data. Next, at operation 314, a message is propagated from the source software module to the destination software module. After the message has been propagated, the thread-count for the appropriate processor is updated (e.g., decremented from one to zero). Following either one of the operations 316 or 310, the data management method 300 ends.

As will be appreciated by those skilled in the art, the message is propagated at operation 314 without preventing other threads from propagating data. More particularly, several threads running on different processors can concurrently propagate data between the software modules when operation 302 determines that data can be propagated concurrently. It should be noted that the propagation of the message at operation 314 can be achieved without having to acquire either the data flow lock or the synchronization queue lock described respectively with respect to operations 304 and 306. As a result, there is no need for a thread to wait until a lock can be acquired before a message can be propagated to a subsequent software module. Thus, multiple threads running on different processors can concurrently propagate data.

Figs. 4A and 4B illustrate a data flow management method 400 configured according to one embodiment of the invention. The data flow management method 400 manages flow of messages between software modules configured in a stack. The data flow management 400 can be implemented to utilize a segregated synchronized queue to facilitate temporary storage, processing and passing of messages by one software module to another software module. In this embodiment, it is assumed that the segregated synchronization queue has an event queue container and a plurality of data queue containers. Additional details about segregated synchronization queues can be found in U.S. Patent Application No. \_\_\_\_\_ (Att.Dkt.No.: SUN1P397/P4611), entitled "METHOD AND APPARATUS FOR MANAGING DATA PROPAGATION BETWEEN SOFTWARE MODULES" which has been incorporated by reference above.

Initially, at operation 402, a determination is made as to whether any operational events are being processed. If it is determined at operation 402 that no operational events are being processed, a determination is made at operation 404 as to whether any operational events are pending to be processed. Typically, messages should not be processed and/or passed to the next software module when an operational event is being processed or is pending to be processed. The determination at operation 404 can efficiently be performed by

checking the event queue container of the segregated synchronized queue to determine whether the event queue container contains a message pertaining to an operational event (event). An event may effect the flow of data and/or how data is to be processed. For example, one such operational event may be related to changing the path of data flow to facilitate re-routing messages through a different physical router. Typically, data pertaining to operational events should be processed before data not pertaining to events is processed. If no events are pending or being processing, the data flow management method 400 proceeds to operation 406 where a determination is made as to whether the thread-count for the processor which is running the thread seeking to propagate data is zero. A zero thread-count indicates that no threads running on the processor are currently propagating data between the software modules.

However, a non-zero thread count indicates that a thread running on the processor is propagating data between the software modules. In the case when threat count is not zero and/or if it determined that an event is running at operation 402 and/or it determined that an event is pending, the data flow management method 400 proceeds to operation 408 where a request is made to acquire a data flow lock. The data flow lock can be implemented as an exclusive software lock that excludes any thread except the thread that has acquired the data flow lock from propagating data between the software modules. When the data flow lock becomes available, the data flow management method 400 proceeds to operation 410 where the data flow lock is acquired. It should be noted that when the data flow lock is acquired, the flow of data is blocked for any other threads that need to propagate data between the software modules. With other threads being prevented from propagating data, the data flow management method 400 proceeds to operation 412 where a determination is made as to whether any threads are accessing the synchronization queue of the software module where messages are being propagated to (destination software module). This determination can be made by checking the synchronization queue counter of the destination software module. If the operation 412 determines that a thread is accessing the synchronization queue of the destination software module, the data flow

management method 400 proceeds to operation 414 where a request is made to acquire the synchronization for the synchronization queue lock of the destination software module. When the synchronization queue lock becomes available, the data flow management method 400 proceeds to operation 416  
5 where the lock for the synchronization queue of the destination software module is acquired. When the synchronization queue lock has been acquired, a message can be placed at the end of the synchronization queue of the destination software module at operation 418. Here, the messages is moved from the source software module to the destination software module. Next, at  
10 operation 420, both the data flow lock (acquired at operation 410) and the synchronization queue lock of the destination software module (acquired at operation 416) are released. Following the operation 420, the data flow management method 400 ends.

On the other hand, if operation 406 determines that the thread-count for a  
15 the appropriate processor is zero, the data flow management method 400 proceeds to operation 424 where the thread-count for the processor is incremented by one. It should be noted that when operation 412 determines that there are no threads accessing the synchronization queue of the destination software module (i.e., synchronization queue count = 0), the data flow  
20 management method 400 also proceeds to operation 424 after the data flow lock is released at operation 422. In both cases, following the incrementing of the thread-count for the appropriate processor, the data flow management method 400 proceeds to operation 426 where a determination is made as to whether the synchronization queue of the destination software module is empty. It should be  
25 noted that if the synchronization queue contains a message, the message from the source software module is not directly propagated to the (main) queue of the destination software module. Rather, the message from the source software module is placed in the synchronization queue of the destination software module. Accordingly, if the operation 426 determines that the synchronization  
30 queue is not empty, the synchronization queue counter of the destination software module is incremented by one at operation 428. At operation 430 a message is retrieved from the head of the synchronization queue of the

destination software module and propagated to the (main) queue of the destination software module. Next, at operation 432, the message from the source module is placed at the bottom of the synchronization queue of the destination software module. The synchronization queue counter of the destination software module is decremented by one at operation 434. The thread-count is then decremented by one at operation 436. Following the operation 436, the data flow management method 400 ends.

Alternatively, when the determination at operation 426 determines that the synchronization queue of the destination software module is empty, the data flow management method 400 proceeds to operation 438 where the message from the source software module is propagated to the (main) queue of the destination software module. In this case, there is no earlier message waiting in the synchronization queue of the destination software module so that the message can be propagated directly to the (main) queue. Following the operation 436 where the thread-count is decremented by one, the data flow management method 400 ends.

It should be noted that when operations 424-438 are performed the message propagation is achieved without using any software locks. The thread-count allows one thread from each one or more processors to utilize such lock-free data propagation. The synchronization queue count is used also to permit lock-free data propagation but only when no other threads are accessing the synchronization queue. The result of lock-free data propagation is significant improvement in data propagation. Data propagation is achieved more efficiently because less blocking by software locks occurs and because concurrent access to synchronization queue is permitted.

Figs. 5A and 5B illustrate a computing environment 500 in which a representative message M4 is propagated from a (main) queue 501 of a source software module 502 to a destination software module 504 in accordance with one situation of the present invention. As shown in Fig. 5A, queue 506 and segregated synchronization queue 508 are provided for the destination software



module 504. In addition, a propagation controller 509, a thread-count 510 for the appropriate processor and a synchronization queue count 512 are provided to facilitate the propagation of the message M4 from the queue 501 of the source software module 502 to the destination software module 504. In the environment depicted in Fig. 5A, both the thread-count 510 and synchronization queue count 512 are at zero. A zero thread count indicates that there is no thread propagating data on the appropriate processor. Similarly, a zero on the synchronization queue counter 512 indicates that there is no thread accessing the synchronization queue 512. It should be noted that in the situation depicted in Fig. 5A, the synchronization queue 508 is empty. That is, there are no events pending in the event queue 514 and there are no data messages in the data queues 516. In this situation, similar to the case of operation 438 of Fig. 4, the thread-count 510 is incremented to 1 and the message M4 is propagated to the queue 506 of the destination software module 504. Fig. 5B illustrates the computing environment 500 after M4 is propagated. It should be noted that the thread-count 510 is decremented back to zero following the propagation of message M4 to the queue 506 (not shown in Fig. 5B).

Figs. 5C and 5D illustrate propagation of a message M4 from the (main) queue 501 of the source software module 502 to a destination software module 504 in accordance with another situation of the present invention. Fig. 5C depicts an environment similar to Figs. 5A and 5B. However, it should be noted that in the situation depicted in Fig. 5C, the synchronization queue 508 is not empty. More specifically, the data queue 516 of the synchronization data queue 508 contains data messages MA and MB. In this situation, the message M4 is placed at the bottom of the data queue 516 and the message MA is retrieved from the top of the data queue 506 and propagated to the queue 506 of the destination software module 504. (operations 430 and 432 of Fig. 4).

Figs. 6A and 6B illustrate an example of a computer system that may be used in accordance with the invention. Fig. 6A shows a computer system 601 that includes a display 603, screen 605, cabinet 607, keyboard 609, and mouse 611. Mouse 611 may have one or more buttons for interacting with a graphical

user interface. Cabinet 607 houses a CD-ROM drive 613, system memory and a hard drive (see Fig. 6B) which may be utilized to store and retrieve software programs incorporating computer code that implements the invention, data for use with the invention, and the like. Although CD-ROM 615 is shown as an exemplary computer readable storage medium, other computer readable storage media including floppy disk, tape, flash memory, system memory, and hard drive may be utilized. Additionally, a data signal embodied in a carrier wave (e.g., in a network including the Internet) may be the computer readable storage medium. In one implementation, an operating system for the computer system 601 is provided in the system memory, the hard drive, the CD-ROM 615 or other computer readable storage medium and serves to incorporate the computer code that implements the invention.

Fig. 6B shows a system block diagram of computer system 601 used to execute the software of an embodiment of the invention. As in Fig. 6B, computer system 601 includes monitor 603 and keyboard 609, and mouse 611. Computer system 601 further includes subsystems such as a central processor 651, system memory 653, fixed storage 655 (e.g., hard drive), removable storage 657 (e.g., CD-ROM drive), display adapter 659, sound card 661, speakers 663, and network interface 665. The central processor 651, for example, can execute computer program code (e.g., an operating system) to implement the invention. An operating system is normally, but necessarily) resident in the system memory 653 during its execution. Other computer systems suitable for use with the invention may include additional or fewer subsystems. For example, another computer system could include more than one processor 651 (i.e., a multi-processor system) or a cache memory.

The system bus architecture of computer system 601 is represented by arrows 667. However, these arrows are illustrative of any interconnection scheme serving to link the subsystems. For example, a local bus could be utilized to connect the central processor to the system memory and display adapter. Computer system 601 shown in Fig. 6B is but an example of a computer system suitable for use with the invention. Other computer architectures having different configurations of subsystems may also be utilized.

The invention can use a combination of hardware and software components. The software can be embodied as computer readable code (or computer program code) on a computer readable medium. The computer readable medium is any data storage device that can store data which can be  
5 thereafter be read by a computer system. Examples of the computer readable medium include read-only memory, random-access memory, CD-ROMs, magnetic tape, optical data storage devices. The computer readable medium can also be distributed over a network coupled computer systems so that the computer readable code is stored and executed in a distributed fashion.

10 The advantages of the invention are numerous. Different embodiments or implementations may have one or more of the following advantages. One advantage of the invention is that data can be propagated concurrently between software modules by a plurality of threads. Another advantage of the invention is that lock-free data propagation is made possible. Yet another advantage of the  
15 invention is that propagation of data can be achieved more quickly and more efficiently.

The many features and advantages of the present invention are apparent from the written description, and thus, it is intended by the appended claims to cover all such features and advantages of the invention. Further, since numerous  
20 modifications and changes will readily occur to those skilled in the art, it is not desired to limit the invention to the exact construction and operation as illustrated and described. Hence, all suitable modifications and equivalents may be resorted to as falling within the scope of the invention.

*What is claimed is:*

## CLAIMS

1. A method for managing flow of messages between two software modules, said method comprising:

(a) determining whether a first message can be propagated by a first thread running on a first processor between the two software modules while allowing a second thread running on a second processor to propagate a second message between the two software modules; and

(b) propagating the first message between the two software modules while allowing the second thread to propagate the second message between the two software modules when said determining (a) determines that the first message can be propagated by the first thread while allowing the second thread to propagate the second message between the two software modules.

2. A method as recited in claim 1, wherein the first and second threads concurrently propagate respective portions of the first and second messages between the two software modules.

3. A method as recited in claim 1, wherein the method further comprises:

(c) blocking the second thread from propagating the second message between the two software modules when said determining (a) determines that the first message cannot be propagated by the first thread while allowing the second message to propagate between the two layer software modules.

4. A method as recited in claim 2, wherein said blocking (c) of the second thread from propagating the message between the two software modules is achieved by providing a lock which can be acquired by the first thread.

5. A method as recited in claim 1, wherein the method further comprises:

setting a first indicator for the first processor to indicate that the first processor is propagating when said determining (a) determines that the first

message can be propagated by the first thread while allowing the second thread to propagate the second message between the two layer software modules.

setting the first indicator for the first processor to indicate that the first processor is not propagating when said propagating (b) has propagated the first message between the two software modules.

6. A method as recited in claim 1, wherein said determining (a) comprises:

(a1) determining whether an event is being processed or is pending to be processed; and

(a2) determining whether a thread-count for the first processor is zero,

and

wherein said determining (a) determines that the first thread can propagate the first message without blocking the second thread from propagating the second message when said determining (a) determines that no events is being processed or pending and said determining (a) determines that the thread-count for the first processor is zero.

7. A method as recited in claim 1, wherein the method further comprises:

(a1) determining whether a synchronization queue associated with one of the software modules contains one or more data messages.

8. A method as recited in claim 6, wherein the method further comprises:

propagating the first message to the bottom of the synchronization queue when the determining (a1) determines that the synchronization queue contains one or more messages; and

propagating a message from the head of the synchronization queue when the determining determines (a1) that the synchronization queue contains one or more messages.

9. A method as recited in claim 6, wherein the method further comprises:

propagating the first message to the bottom of a queue of one of the software modules when the determining determines (a1) that the synchronization queue does not contain one or more messages.

5 10. A method as recited in claim 1, wherein the two software modules are implemented in a stack as STREAMS modules.

11. A computer system comprising:

a plurality of processors;

10 first and second software modules, the second software module having a main queue suitable for storing messages and an auxiliary queue suitable for storing messages that are not stored in the main queue; and

a propagation controller operating to enable at least two processors of said plurality of processors to concurrently propagate messages between the first and the second software modules.

15 12. A computer system as recited in claim 11,

wherein the propagation controller comprises:

a thread-count for one of said plurality of processors; and

a queue count for the auxiliary queue.

13. A computer system as recited in claim 12,

20 wherein the auxiliary queue is a synchronization queue and the queue count is a synchronization queue count.

14. A computer system as recited in claim 11, wherein the at least two processors of said plurality of processors concurrently propagate a message from the first software module to the auxiliary queue of the second software module.

15. A computer system as recited in claim 11, wherein the at least two processors of said plurality of processors concurrently propagate a message from the first software module to the main queue of the second software module.

16. A computer system as recited in claim 16, wherein the first and second software modules are implemented in a stack as STREAMS modules.

17. A computer readable media including computer program code for managing flow of messages between two software modules, said computer readable media comprising:

computer program code for determining whether a first message can be propagated by a first thread running on a first processor between the two software modules while allowing a second thread running on a second processor to propagate a second message between the two software modules; and

computer program code for propagating the first message between the two software modules while allowing the second thread to propagate the second message between the two software modules when said computer program code for determining determines that the first message can be propagated by the first thread while allowing the second thread to propagate the second message between the two software modules.

18. A computer readable media as recited in claim 17, wherein the first and second threads concurrently propagate respective portions of the first and second messages between the two software modules.

19. A computer readable media as recited in claim 18, wherein the computer readable media further comprises:

computer program code for setting a first indicator for the first processor to indicate that the first processor is propagating when said determining (a) determines that the first message can be propagated by the first thread while

allowing the second thread to propagate the second message between the two layer software modules.

computer program code for setting the first indicator for the first processor to indicate that the first processor is not propagating when said propagating (b)  
5 has propagated the first message between the two software modules.

20. A computer readable media as recited in claim 19, wherein the two software modules are implemented in a stack as STREAMS modules.




# METHOD AND APPARATUS FOR CONCURRENT PROPAGATION OF DATA BETWEEN SOFTWARE MODULES

5

## ABSTRACT OF THE DISCLOSURE

Improved techniques for managing propagation of data through software modules used by computer systems are disclosed. The improved techniques allow a plurality of threads running on different processors to propagate data between software modules by providing concurrent access to queued data by multiple threads. The techniques are particularly well suited for use with the STREAMS model that uses software models arranged in a stack to provide data communications.

10



**Fig. 1**


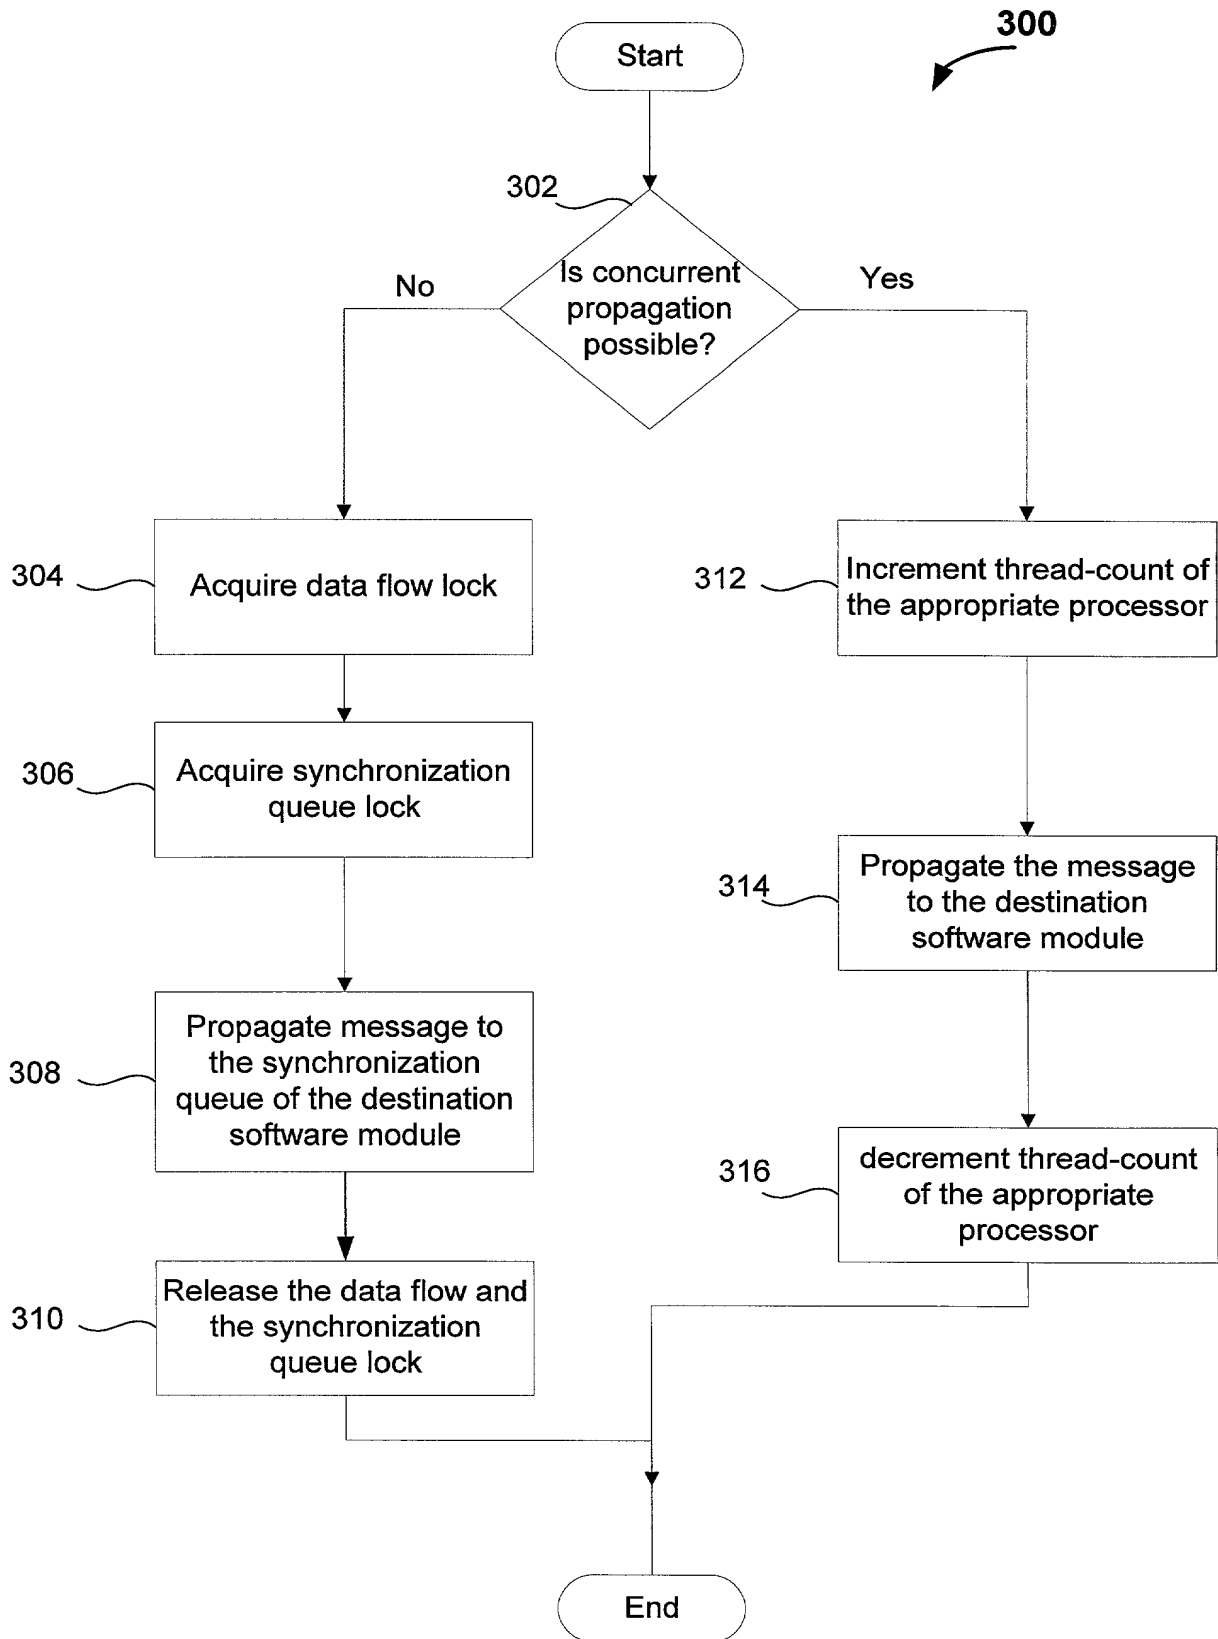


Fig. 2



**Fig. 3**

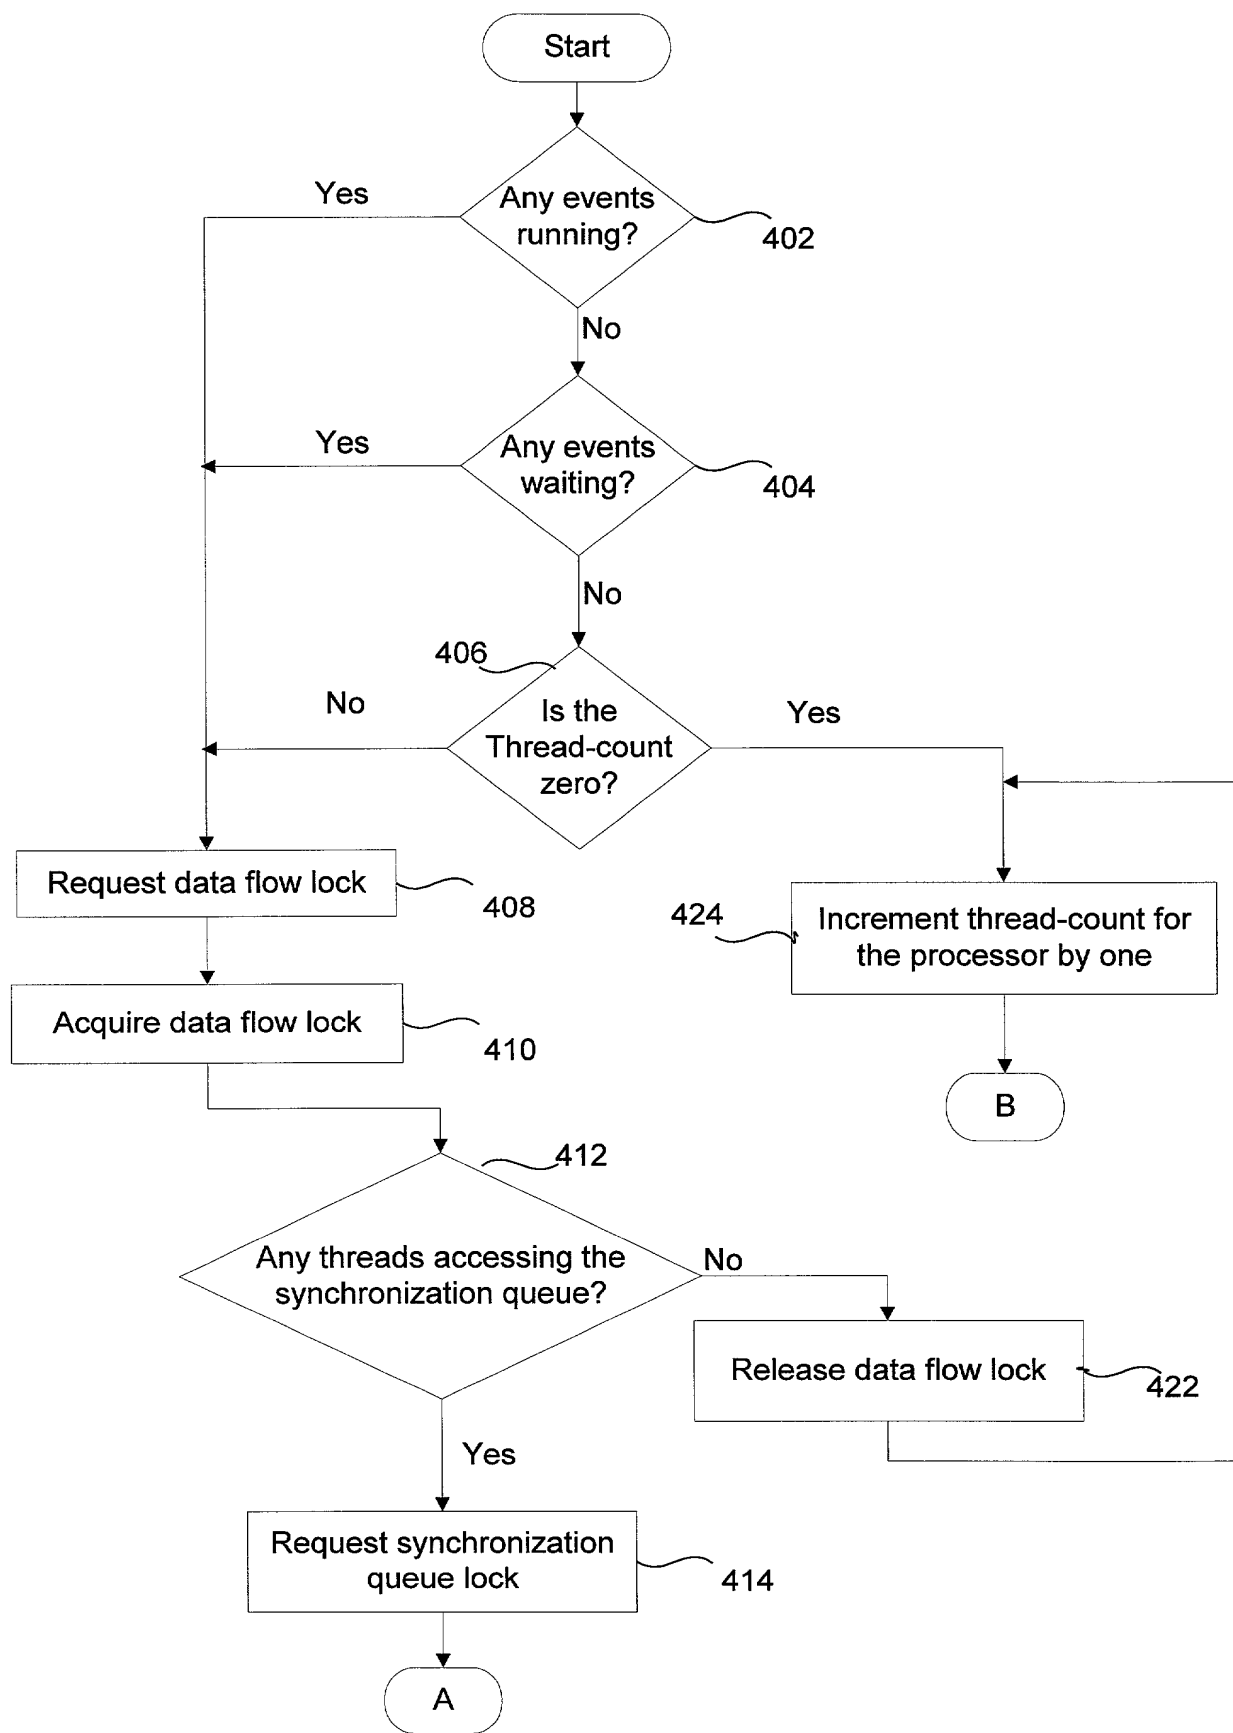
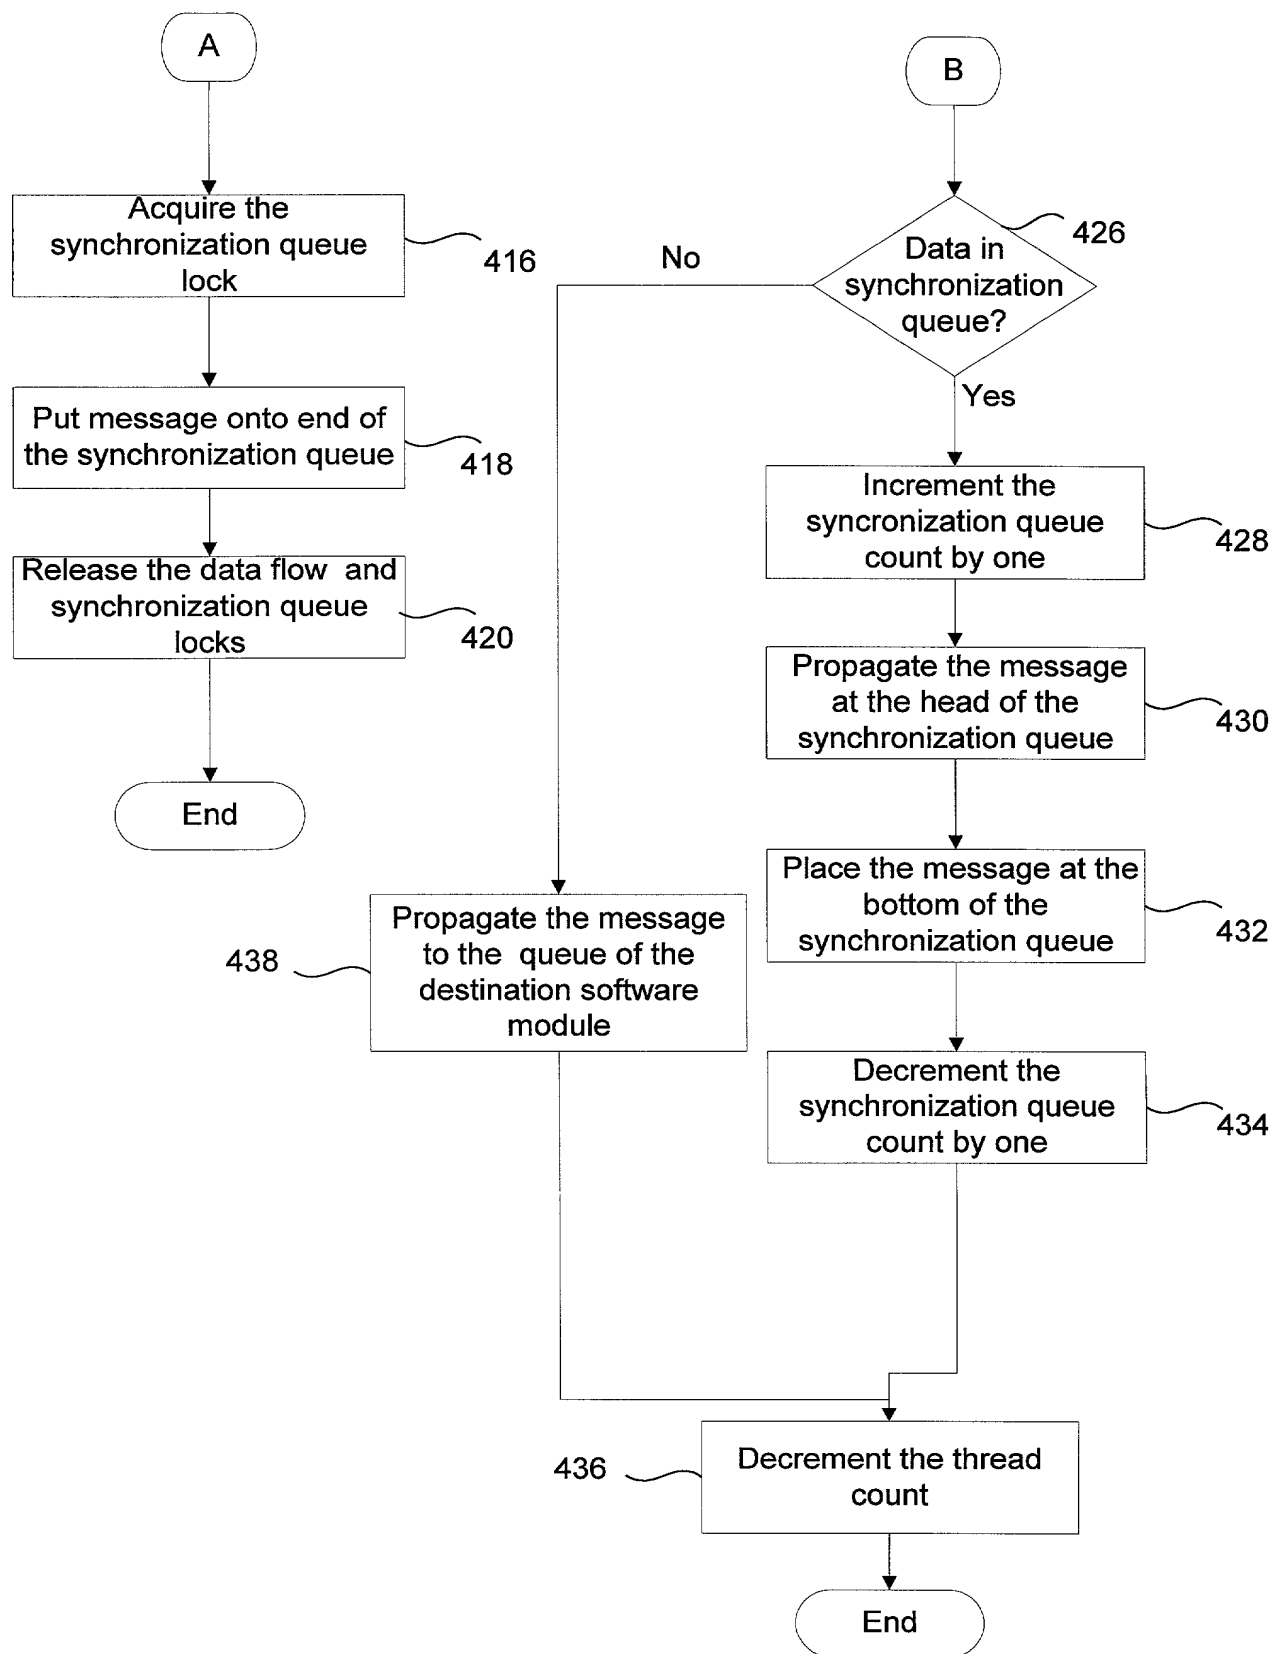



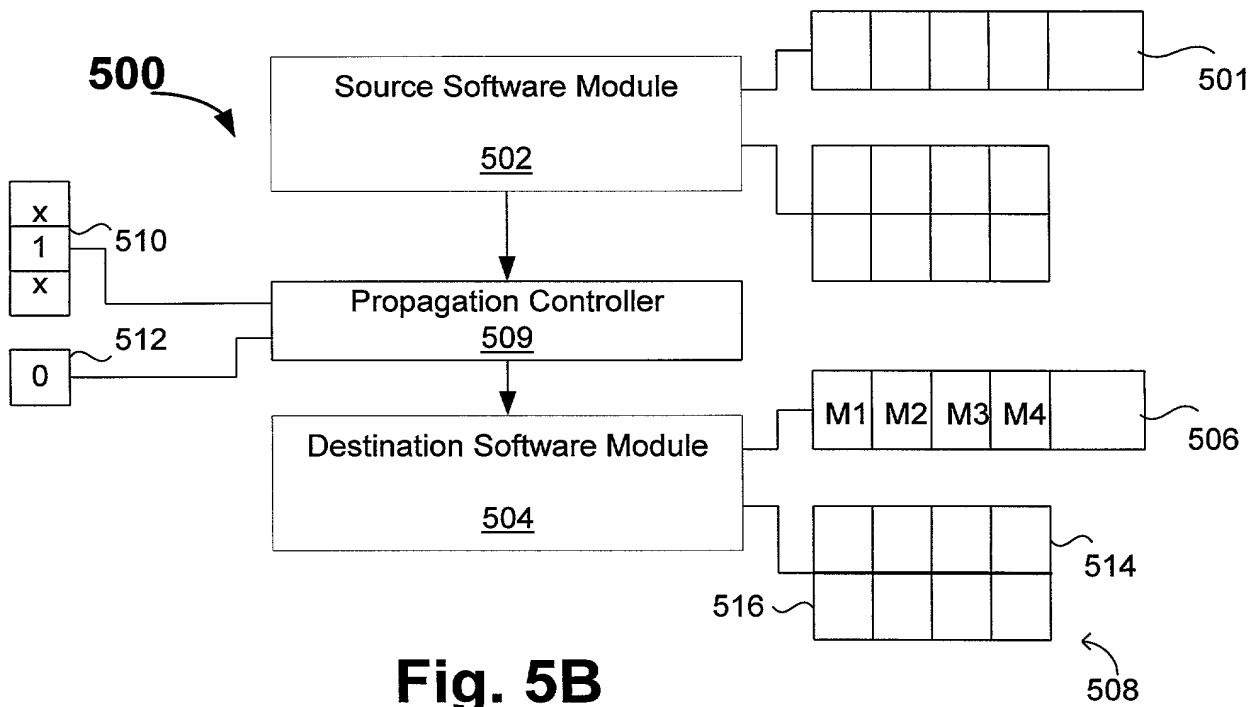
Fig. 4A





**Fig. 4B**




**Fig. 5A**



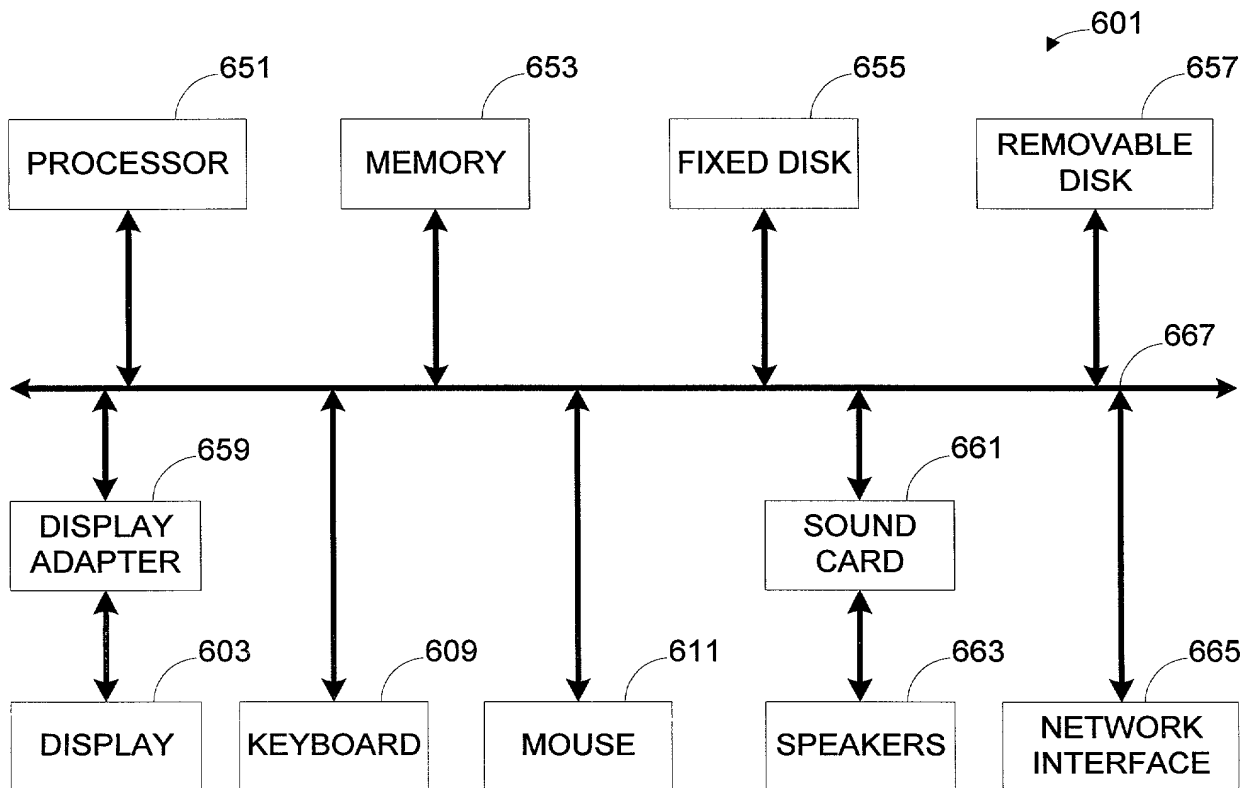
**Fig. 5B**







**FIG. 6A**



**FIG. 6B**

# DECLARATION AND POWER OF ATTORNEY FOR ORIGINAL U.S. PATENT APPLICATION

Attorney's Docket No. SUN1P398/P4686

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: METHOD AND APPARATUS FOR CONCURRENT PROPAGATION OF DATA BETWEEN SOFTWARE MODULES the specification of which,

(check one)

1. ☒ is attached hereto.
2. ☐ was filed on \_\_\_\_\_ as  
U.S. Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.
3. ☐ was filed on \_\_\_\_\_ as  
International PCT Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

## Prior Foreign Application(s)

I hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

			Priority Benefits Claimed?
			Yes ___ No ___
_____ (Application No.)	_____ (Country)	_____ (Filing Date)	
_____ (Application No.)	_____ (Country)	_____ (Filing Date)	Yes ___ No ___

## Provisional Application(s)

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below:

_____ (Application No.)	_____ (Filing Date)
_____ (Application No.)	_____ (Filing Date)

### Prior U.S. Application(s)

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:


_____ (Application No.)	_____ (Filing Date)	_____ (Status - patented, pending, abandoned)
_____ (Application No.)	_____ (Filing Date)	_____ (Status - patented, pending, abandoned)

### Power of Attorney

And I hereby appoint the law firm of **Beyer Weaver Thomas & Nguyen, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

### Direct Correspondence To:

**Customer Number: 022434**  
**BEYER WEAVER THOMAS & NGUYEN, LLP**  
P.O. Box 130  
Mountain View, CA 94042-0130



jc525 U.S. PTO  
09/513706  
02/25/00

### Direct Telephone Calls To:

**C. Douglass Thomas at telephone number (650) 961-8300**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

### Typewritten Full Name of

Sole or First Inventor: Paul F. Lodrige Citizenship: US

Inventor's signature: \_\_\_\_\_ Date of Signature: \_\_\_\_\_

Residence: (City) Palo Alto (State/Country) CA/USA

Post Office Address: 784 Holly Oak Drive, Palo Alto, CA 94303

Second Inventor: Randy S. Fishel Citizenship: US

Inventor's signature: \_\_\_\_\_ Date of Signature: \_\_\_\_\_

Residence: (City) Sunnyvale (State/Country) CA/USA

Post Office Address: 854 Laburnum, Sunnyvale, CA 94086